

## LA-UR-18-27482

Approved for public release; distribution is unlimited.

Title: Job Profiling and Cluster Monitoring for Users

Author(s): Duarte, Marco Antonio  
Rindels, Raymond Luke  
Everson, Kody Jacob  
Lopatina, Olena  
Yilk, Todd Alan  
Lopez, Anthony Louis

Intended for: HPC mini showcase, ISTI Day

Issued: 2018-08-06

---

**Disclaimer:**

Los Alamos National Laboratory, an affirmative action/equal opportunity employer, is operated by the Los Alamos National Security, LLC for the National Nuclear Security Administration of the U.S. Department of Energy under contract DE-AC52-06NA25396. By approving this article, the publisher recognizes that the U.S. Government retains nonexclusive, royalty-free license to publish or reproduce the published form of this contribution, or to allow others to do so, for U.S. Government purposes. Los Alamos National Laboratory requests that the publisher identify this article as work performed under the auspices of the U.S. Department of Energy. Los Alamos National Laboratory strongly supports academic freedom and a researcher's right to publish; as an institution, however, the Laboratory does not endorse the viewpoint of a publication or guarantee its technical correctness.

# Job Profiling and Cluster Monitoring for Users

Marco Duarte, Kody Everson, Raymond Rindels

CSCNSI 2018





# Introduction

- Users have no way of seeing how their jobs perform at a hardware level
- Metrics are scattered across different services
- Slurm scheduling could be more resource aware
- Variability in shared resource usage can cause variability in performance



# Solutions

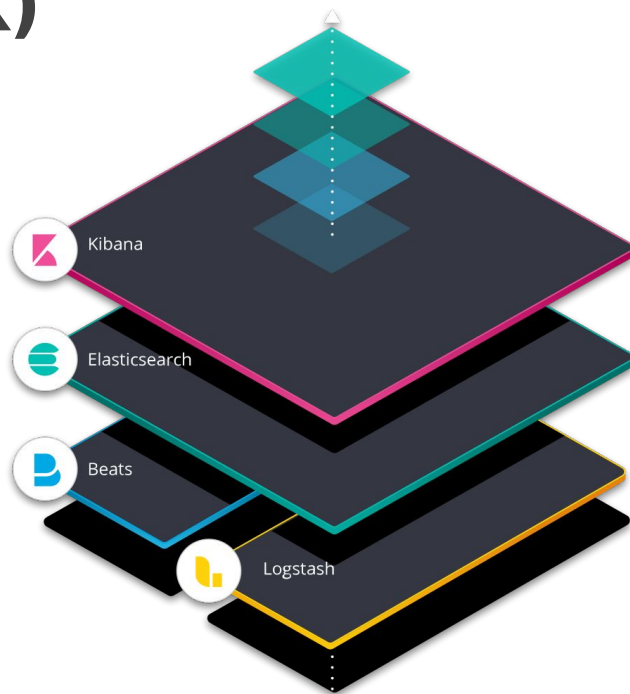
- Aggregate all the metrics into a centralized database
- Personalized dashboards to show users what will benefit them
- Use data gathered during Slurm jobs to create job profiles

# The Software Pipeline



# The Elastic Stack (ELK)

- Set of open source projects
- Collects and visualizes data from many different sources
- Each component can be used independently with other tools





# Elasticsearch

- Data storage
- Can distribute load and data across its own cluster
- Scales by adding more nodes
- Extremely fast







# Logstash


- Data processing
- Ingests data from many different sources
- Makes it easy to parse things like logs
- Filters out unnecessary data



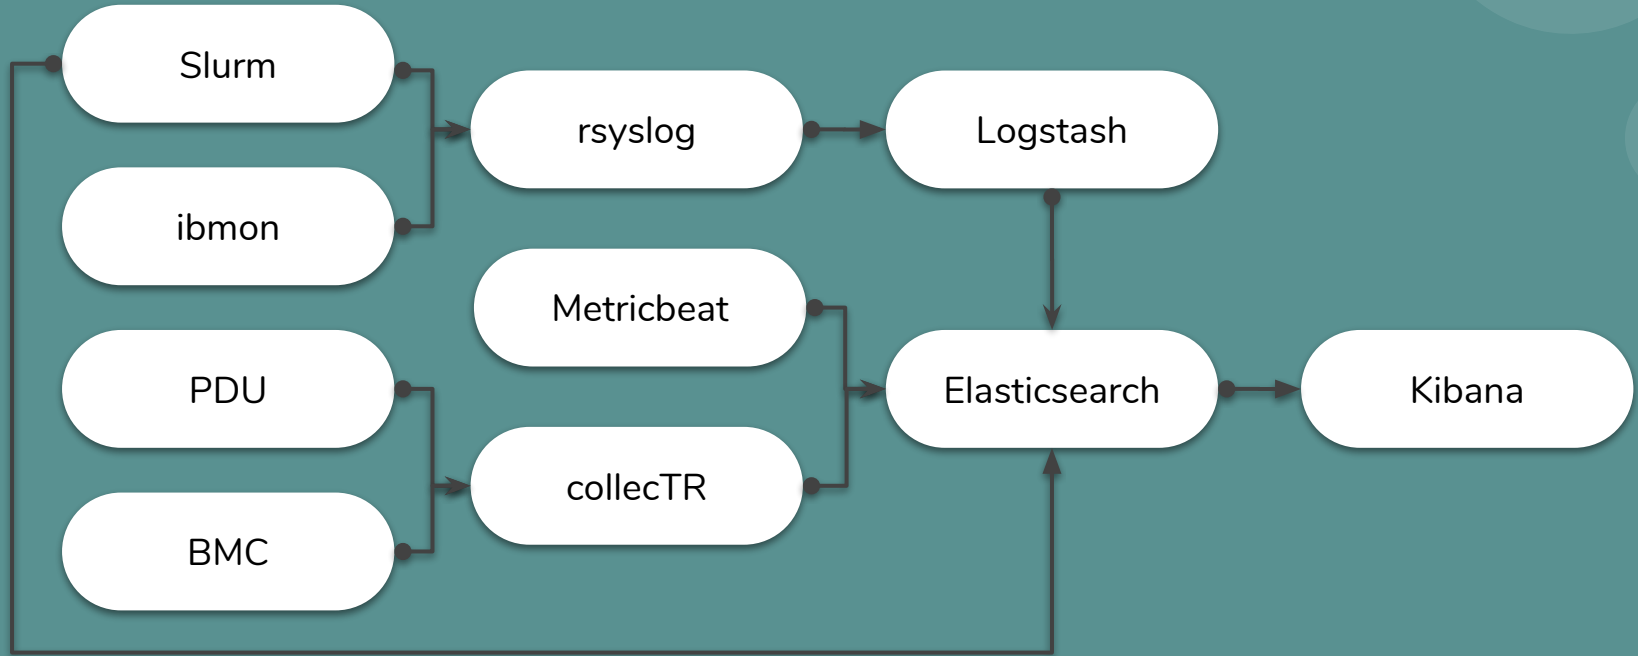


# Kibana

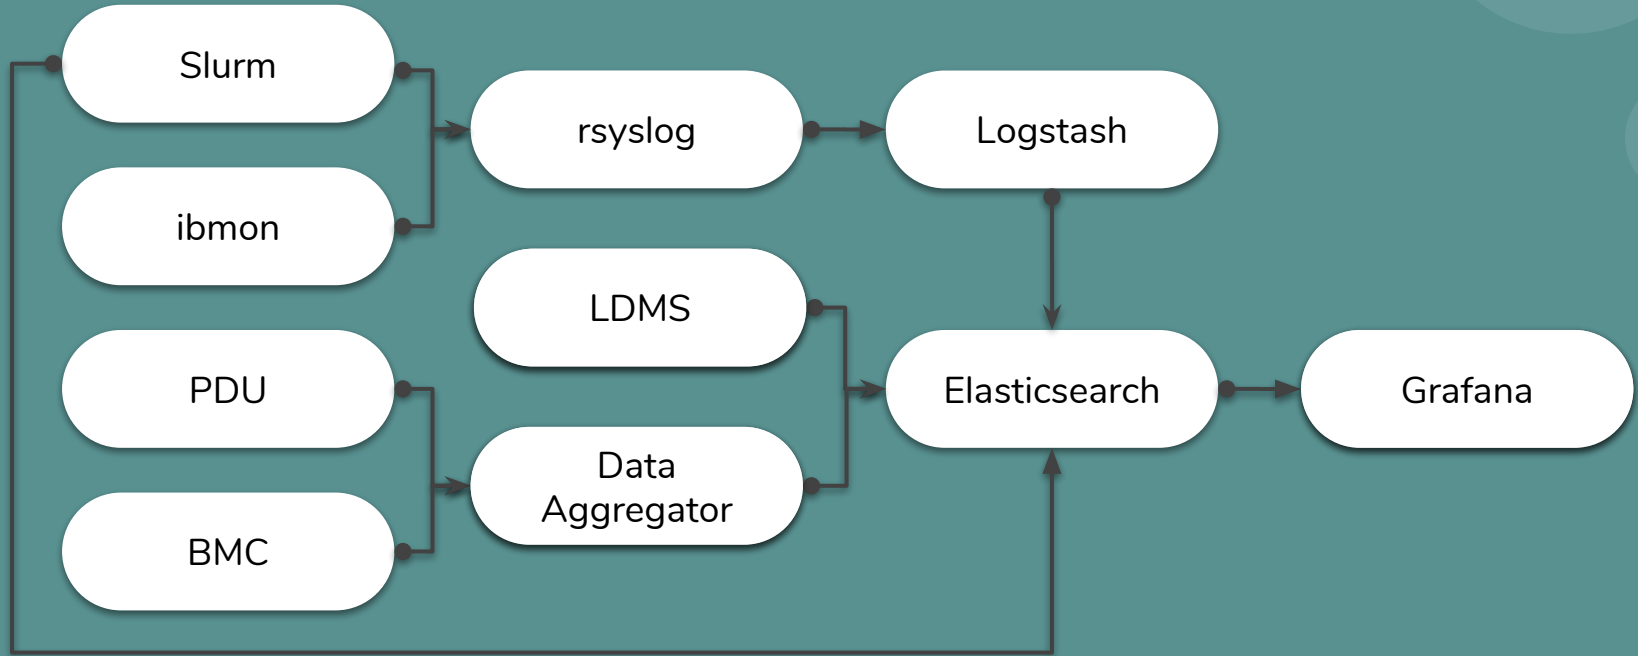
- Data visualization
- Easy-to-use interface
- Includes templates as well as custom tools
- Paid version (XPack) includes PDF reports and alerting



# Our Setup



# Potential LANL Setup





# Pros and Cons

+

- Fine-grain control
- Open source
- Scalable
- Modular
- Plugins

-

- Learning curve
- Requires upkeep at scale
- Documentation for recent versions is lacking


# Visualizing Data





# Dashboards

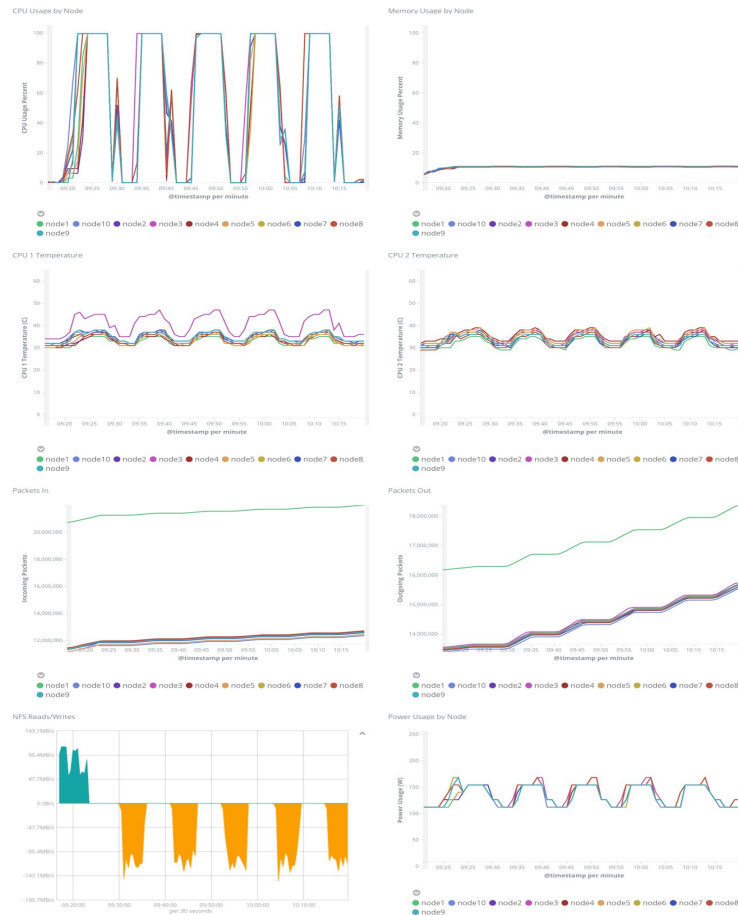
- Some dashboards need to be constantly updated, others need to be over a fixed interval
- Keep cluster scale in mind





# Types of Dashboards

- System Overview
- Node Overview
- User Overview
- Job Overview
- Job Profile






# System Overview

Job Overview | User Overview

11  
Nodes


Average Inlet Temperatures



Average Exhaust Temperatures




Average CPU Temperatures




Infiniband Congestion



Shared Storage Space




Average CPU Usage




Top Hosts By CPU

head.teamred.com	2.74%
node9	0.17%
node8	0.13%
node6	0.12%
node7	0.12%
node3	0.12%
node5	0.12%
node10	0.12%
node2	0.12%
node1	0.11%
node4	0.11%

System Power (W)



Average Memory Usage




# System Overview

Job Overview | User Overview

# 11

Nodes


Average Inlet Temperatures



Average Exhaust Temperatures




Average CPU Temperatures




Infiniband Congestion



Shared Storage Space




Average CPU Usage




Top Hosts By CPU

node4		100%
node7		100%
node2		100%
node1		100%
node8		100%
node5		100%
node3		100%
node9		99.98%
node10		99.97%
head.teamred.com		2.08%
node6		0.15%

System Power (W)



Average Memory Usage






# Node View

[System Overview](#) | [Job Overview](#) | [User Overview](#)

**node7**  
Node Name

## Temperature



## Power


Power Usage (W)






## CPU


CPU Usage



▼ user 99.68% system 0.08% nice 0% irq 0.23% softirq 0.02% iowait 0%




System Load




▼ 1m 31.86 5m 21.73 15m 15.55

## Memory

Memory Usage



▼ Cache -318.1MB Free 59.6GB Used 3.5GB



Memory usage

**3.5GB**

Total Memory 62.8GB

## Top Processes By CPU


metricbeat	5%
rcuos/1	0.2%
kworker/24:0	0%
kworker/1:2	0%
kworker/19:0	0%
kworker/17:1	0%
kworker/0:1H	0%
kworker/0:0	0%
dbus-daemon	0%
RealTest	0%

## Top Processes By Memory


systemd-journal	0.1%
metricbeat	0.1%
kworker/17:1	0%
kworker/0:1H	0%
kworker/0:0H	0%
kworker/0:0	0%
ksoftirqd/7	0%
ksoftirqd/0	0%
dbus-daemon	0%
RealTest	0%

## Network

### Network Traffic (Packets)



### Network Traffic (Bytes)



### Interfaces by Incoming Traffic

Inbound Traffic

**1.2KB/s**

kibana x Transferred 10.1MB

em1	77.9GB
lo	129.9MB
ib0	94.4MB
eno4	0.0B
eno3	0.0B

### Interfaces by Outgoing traffic


em1	127.9GB
ib0	7.8GB
lo	129.9MB
eno4	0.0B
eno3	0.0B

Outbound Traffic

**7.1KB/s**

Total Transferred 7.0MB

## Cluster Heatmap



### Queued jobs

Job ID ↕	Submit Time ↕	Start Time ↕	Nodes ↕	State ↕	Count ↕
578	July 25th 2018, 09:32:31.441	-	-	PENDING	1
577	July 25th 2018, 09:32:11.519	-	-	PENDING	1
576	July 25th 2018, 09:32:10.673	July 25th 2018, 09:32:17.993	node[7-10]	STARTED	1
571	July 25th 2018, 09:31:57.297	July 25th 2018, 09:31:58.102	node[5-6]	STARTED	1

Export: Raw  Formatted 

# User Overview

System Overview | Job Overview

Average Wait in Queue (Mins)

0.18

3

Running Jobs


2

Pending Jobs


Top 5 Slurm Users by Job Count




System Power (W)



NFS Reads/Writes



Packets Incoming



Packets Outgoing





# Scaling

- This approach works great for a 10 node cluster, but scaling up is a potential issue
- Dashboards could be redesigned to be more intuitive at a large scale





# Slurm



## Job Information

- Completed jobs use the Elasticsearch plugin to send data to Elasticsearch
- Pending and running jobs need to be picked up from rsyslog through Logstash





# Profiling

- Matching up system metrics with job start and end times allows us to see how a job runs
- This data can help users optimize jobs and understand variability in performance



# Job Overview

System Overview | User Overview

264

Slurm Jobs on Record

Average Wait in Queue (Mins)

2.09

Average Run Time (Mins)

3.93

Job State

Select...

User

Select...

Job ID Range

188

546

Clear form

Cancel changes

Apply changes

Slurm Job Table [Completed]

Slurm Job ID ▾	User ▾	Job Name ▾	Nodes ▾	Submit Time ▾	Start Time ▾	End Time ▾	State ▾	Run Time ▾	Job Link ▾	PDF Link ▾
546	mpiuser	hpl	(null)	July 23rd 2018, 13:56:18.000	July 23rd 2018, 13:58:20.000	July 23rd 2018, 13:58:20.000	CANCELLED	0M 0S	<a href="#">Click for More Info</a>	<a href="#">PDF</a>
545	mpiuser	hpl	(null)	July 23rd 2018, 13:56:17.000	July 23rd 2018, 13:58:20.000	July 23rd 2018, 13:58:20.000	CANCELLED	0M 0S	<a href="#">Click for More Info</a>	<a href="#">PDF</a>
544	mpiuser	hpl	(null)	July 23rd 2018, 13:56:16.000	July 23rd 2018, 13:58:20.000	July 23rd 2018, 13:58:20.000	CANCELLED	0M 0S	<a href="#">Click for More Info</a>	<a href="#">PDF</a>
543	mpiuser	hpl	(null)	July 23rd 2018, 13:56:15.000	July 23rd 2018, 13:58:20.000	July 23rd 2018, 13:58:20.000	CANCELLED	0M 0S	<a href="#">Click for More Info</a>	<a href="#">PDF</a>
542	mpiuser	hpl	(null)	July 23rd 2018, 13:56:15.000	July 23rd 2018, 13:58:20.000	July 23rd 2018, 13:58:20.000	CANCELLED	0M 0S	<a href="#">Click for More Info</a>	<a href="#">PDF</a>
541	everson	ENJ	node[1-10]	July 23rd 2018, 13:53:32.000	July 23rd 2018, 13:53:33.000	July 23rd 2018, 14:02:04.000	COMPLETED	8M 31S	<a href="#">Click for More Info</a>	<a href="#">PDF</a>
540	everson	ENJ	node[1-10]	July 23rd 2018, 13:44:24.000	July 23rd 2018, 13:44:25.000	July 23rd 2018, 13:50:05.000	CANCELLED	5M 40S	<a href="#">Click for More Info</a>	<a href="#">PDF</a>

# Job View

System Overview | Job Overview | User Overview

**526**  
Slurm Job ID

**everson**

Job Submitted By

**node[1-10]**

Nodes Allocated

**63M 22S**


Run Time

**July 23rd 2018, 09:16:51.000**  
Submit Time


**July 23rd 2018, 09:16:52.000**  
Start Time



**July 23rd 2018, 10:20:14.000**  
End Time

CPU Usage by Node




Memory Usage by Node







Packets In




Packets Out



Power Usage by Node



NFS Reads/Writes




# Future Applications





# Planning for Future Machines

- Deciding on Memory and CPU purchases
- Cooling plans
- Size








# Purchasing Resources


- Some resources are purchased in advance
- Predictive power usage
- Accuracy prevents drawing too much or too little power





# Smart Scheduling

- Resource-aware scheduling
- Shared resource distribution





# Acknowledgements

- Mentors: Lena M Lopatina, Todd Yilk, Anthony Lopez
- CSCNSI: Alfred Torrez, Hunter Easterday, Colette Caskie
- HPC Consultants: Mike Mason, Jesse Martinez, Ben Santos, Joshi Fullop, Brett Kettering

# Questions?

( I'm hungry, time to eat lunch. )

```


o
o  |\_--/|  /  / \ \ / \
    /0 0 \_-- /  \ \ / \
    /      / \ \_ /  \ \ / \
    @_^_@' /  \ \_ /  \ \ / \
    //^_ /  \ \_ /  \ \ / \
    ( // ) |  \ \ / \
    ( / / ) _|_ / ) //
    ( // / ) ' / _ _ / ( ; - .
(( / / )) ,-{      - . | . - ~ .
(( / / / )) ' \ /      ~ . _ . ~
(( /// ))      \ . {      }
(( / ))      .-----~.\      \-'
            ///.-----..>      \
            ///-.-.-.-.-}^-----~

```



# Image Sources

<https://www.elastic.co/elk-stack>



<https://www.elastic.co/products>



<https://www.elastic.co/products>



<https://www.elastic.co/products>



[https://en.wikipedia.org/wiki/Slurm\\_Workload\\_Manager](https://en.wikipedia.org/wiki/Slurm_Workload_Manager)



<https://www.deviantart.com/olivemonkey/art/Slurm-logo-217169884>



<https://officialpsds.com/super-computers-psd-rl1ml>



<http://blockbattle.net/>